

Image Restoration using Deep learning

A Project Report

*Submitted to the APJ Abdul Kalam Technological University
in partial fulfillment of requirements for the award of degree*

in

Master of Computer Application

by

ANILAMOL CHACKO

TKM20MCA-2009

Under the Guidance of

Dr. FOUSIA M SHAMSUDEEN



DEPARTMENT OF MCA

TKM COLLEGE OF ENGINEERING ,KOLLAM

KERALA

July 2022

DEPARTMENT OF MCA
TKM COLLEGE OF ENGINEERING KOLLAM
2020 - 22



CERTIFICATE

This is to certify that the report entitled **Image Restoration using Deep learning** submitted by **ANILAMOL CHACKO** (TKM20MCA-2009), to the APJ Abdul Kalam Technological University in partial fulfillment of the M.C.A degree in Master of Computer Application is a bonafide record of the project work carried out by her under our guidance and supervision. This report in any form has not been submitted to any other University or Institute for any purpose.

(Internal Supervisor)

(Head of the Department)

(External Examiner)

DECLARATION

I ANILAMOL CHACKO hereby declare that the project report **Image Restoration using Deep learning** , submitted for partial fulfillment of the requirements for the award of degree of Master of Computer Application of the APJ Abdul Kalam Technological University, Kerala is a bonafide work done by me under supervision of Dr. FOUSIA M SHAMSUDEEN

This submission represents my ideas in my own words and where ideas or words of others have been included, I have adequately and accurately cited and referenced the original sources.

I also declare that I have adhered to ethics of academic honesty and integrity and have not misrepresented or fabricated any data or idea or fact or source in my submission. I understand that any violation of the above will be a cause for disciplinary action by the institute and/or the University and can also evoke penal action from the sources which have thus not been properly cited or from whom proper permission has not been obtained. This report has not been previously formed the basis for the award of any degree, diploma or similar title of any other University.

Kollam
22-07-2022

ANILAMOL CHACKO

ACKNOWLEDGEMENT

I take this opportunity to express my deepest sense of gratitude and sincere thanks to everyone who helped me to complete this work successfully. I express my sincere thanks to **Dr. FOUSIA M SHAMSUDEEN**, Head of Department, MCA, TKM COLLEGE OF ENGINEERING Kollam for providing me with all the necessary facilities and support.

I would like to place on record my sincere gratitude to my project guide **Dr. FOUSIA M SHAMSUDEEN**, Head of Department, MCA, TKM COLLEGE OF ENGINEERING Kollam for the guidance and mentorship throughout the course.

I profusely thank all other faculty members in the department and all other members of department of MCA, TKM COLLEGE OF ENGINEERING Kollam for their guidance and inspirations throughout my course of study.

I owe thanks to my friends and all others who have directly or indirectly helped me in the successful completion of this seminar.

ANILAMOL CHACKO

ABSTRACT

Recovering the original image by erasing noise and blur is known as image reconstruction or restoration. Contrary to typical restoration tasks, which can be handled by supervised learning, real-world picture deterioration is complicated, and the network is unable to generalise because to the difference in domains between synthetic images and actual old photos. As a result, To offer a unique triplet domain translation network that makes use of both numerous synthetic picture pairings and actual images. We specifically train two Variational autoencoders (VAEs) to convert clean photographs into two latent spaces and aged photos into two latent spaces, respectively. Create a local branch here that targets unstructured faults like noise and blurriness, as well as a global branch with a partly nonlocal obstruct that targets structured flaws like scratches and dust spots. In the latent space, two branches are merged, increasing the ability to fix numerous faults in ancient images. The suggested method for restoring ancient images performs better in terms of visual quality than cutting-edge techniques.

Contents

List of Figures	iv
1 INTRODUCTION	1
2 RELATED WORKS	3
3 METHODOLOGY	6
3.1 Proposed System	6
3.1.1 Latent space translation for restoration	6
3.1.2 Variational autoencoders for image reconstruction	7
3.1.3 Domain arrangement in the latent space of the VAE	9
3.1.4 Latent mapping for restoration	10
3.1.5 Repeatedly restoring deterioration	10
3.1.6 Detection of Defect Regions	11
3.1.7 Face Enhancement	12
3.2 Dataset	14
3.3 Software Requirement	16
4 RESULT AND DISCUSSION	19
4.1 Comparisons	19
4.2 Screenshots	21
5 CONCLUSION	23
6 FUTURE ENHANCEMENTS	25
References	26

List of Figures

3.1	The design of our network for restoration	7
3.2	Architecture view of variational autoencoders	9
3.3	Method of restoration process	10
3.4	Module created particularly as a partly nonlocal block for the latent inpainting	11
3.5	Scratch detection	12
3.6	Face-enhancing progressive generator network	12
3.7	Examples of Face-enhancement	13
3.8	ROC curve for scratch detection	13
3.9	Comparison of quality against state-of-the-art techniques	15
3.10	synthetic photos with scratches	15
3.11	Some defect region detection results on real photos	16
4.1	Comparison of dataset	20
4.2	User study results	21
4.3	Bringing an ancient photograph back to life	21

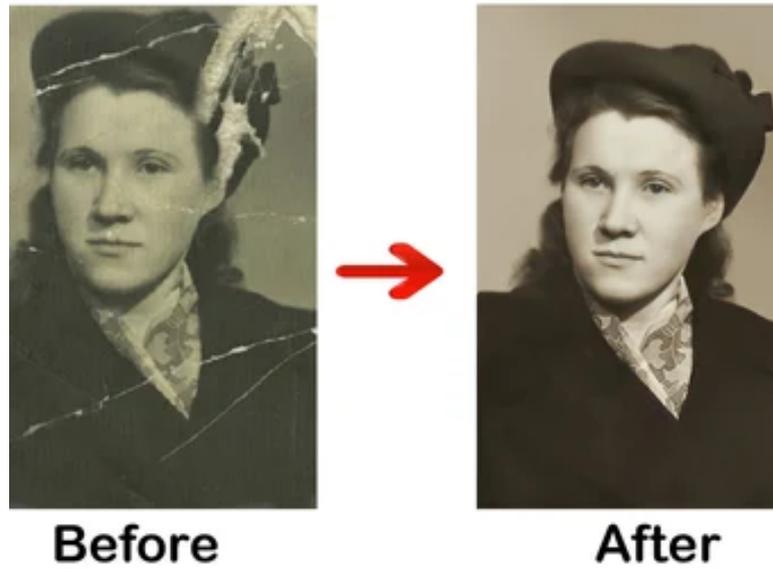
Chapter 1

INTRODUCTION

To capture the joyful moments that would otherwise pass by, photographs are shot. Even if time passes, looking at them might still bring back memories from the past. However, old photo prints decay when stored in unfavourable climatic conditions, irreversibly damaging the priceless visual content. As portable cameras and scanners become more readily available, users may now digitalize the photographs and request restoration from a certified specialist. Manual retouching, on the other hand, is often time-consuming and labor-intensive, rendering a large number of vintage images unsalvageable. As a result, it's intriguing to develop automatic algorithms that can instantly repair old photographs for people who want to resurrect old photos.

Prior to deep learning, many attempts were attempted to repair photos by automatically detecting localised imperfections such as scratches and defects and filling in the damaged parts with inpainting techniques.

In this situation, antique photo restoration is defined as a triplet domain translation problem. Unlike previous image translation algorithms, we employ data from three domains (real old photographs, synthetic images, and associated ground truth) and perform the translation in latent space. Synthetic pictures and genuine photographs are first transferred to the same latent space using a sharing variational autoencoder (VAE). Meanwhile, another VAE is being taught to project clean ground truth pictures into the appropriate latent space. The synthetic picture pairings are then used to learn the



mappings between the two latent spaces, allowing contaminated photographs to be restored to clean ones. Because of the benefit of underlying restoration, the taught latent restoration can generalise well to real-world images.

Chapter 2

RELATED WORKS

Image restoration is the process of calculating the clean, original picture from a corrupted/noisy image. Motion blur, noise, and camera mis-focus are all examples of corruptions. Picture restoration is conducted by correcting the blurring process, which is completed by capturing a point source and restoring the image data lost during the softening process using the point source image, referred as the Point Spread Function (PSF).

Existing picture degradation can be loosely divided into two categories: structural degradation, which includes holes, scratches, and spots, and unstructured degradation, which includes noise, blurriness, colour fading, and low resolution. Traditional works usually impose various image priors on the former unstructured ones, such as non-local self-similarity, atomicity, and global smoothness. Many deep learning-based strategies for dealing with picture degradation, such as denoising, super-resolution, and deblurring, have recently been presented.

Junyuan Xie et al.[1] provide a fresh method for blind inpainting and picture denoising that combines sparse coding with deep neural networks that have already been trained with denoising auto-encoders. They suggest a new DA training approach that enables picture denoise and inpainting inside a single framework. On the straightforward objective of denoising additive white Gaussian noise, our solution performs comparably to conventional linear sparse coding algorithms in the experiments. Furthermore, our non-linear approach successfully solves the, to the best of our

knowledge, unsolved problem of blind inpainting of complicated patterns, which is significantly more difficult. They also demonstrate how the suggested training strategy might enhance DA's performance in unsupervised feature learning tasks.

Julien Mairal et al.[2] suggested a non-local methods and sparse coding techniques to image restoration have been blended into a new image model that presents a unified framework where linked patches are deconstructed using comparable sparsity patterns. According to quantitative and qualitative evaluations using photos contaminated with manufactured or real noise, the suggested technique exceeds the state of art in image demosaicking and denoising tasks. After that, They modify a strategy to tackle other difficult image alteration issues in computational photography, such as deblurring, inpainting, and texture creation in still pictures and video sequences. Our next goal is to incorporate non-uniform noise models into the reconstruction process.

Aneesh G.Nath et al.[3] offers an image super-resolution technique that generates high-resolution pictures directly from compressed data using dictionary learning and non-local patch-based interpolation. The outcomes of the experimental study show that the proposed method can get beyond the drawbacks of other vocabulary learning techniques, which are unable to employ random samples from CS-based image acquisition. Additionally, it effectively solves the neighbourhood issue brought on by single picture super resolution techniques. Visual analysis—both quantitative and subjective—shows how the suggested approach can result in successful outcomes with fewer random samples.

Jian Sun et al.[4] developed a unique nonuniform motion deblurring method based on CNN. From local patches, we develop a useful CNN for estimating motion kernels. They can accurately anticipate the nonuniform motion blur field using an MRF model. This produces cutting-edge motion deblurring effects. They are considering creating a CNN in the future to estimate the general non-uniform blur kernels. Another objective is to develop a CNN system that can estimate and eliminate any non-uniform blurs inside a single framework.

Seungjun Nah et al.[5] proposed a neural network for blind deblurring to estimate

crisp images. Our model, in contrast to earlier research, does not have kernel estimation issues. The suggested model is trained in multi-scale space and uses a coarse-to-fine methodology. Additionally, we created a ground-truth blur dataset that was realistic, allowing for effective supervised learning and thorough testing. According to experimental findings, our technology works better than cutting-edge techniques both qualitatively and quantitatively while being significantly faster.

Chapter 3

METHODOLOGY

3.1 Proposed System

Old photo repair is more difficult than typical image restoration operations. First, there is always a domain gap between genuine and synthetic photographs since ancient photos have far more intricate degradation that is challenging to realistically depict. As a result, by learning exclusively from artificial data, the network typically struggles to generalise well to real photographs. Second, the degradations in ancient photographs are a compound, necessitating diverse solutions for restoration. On the other hand, to guarantee structural consistency, structural faults like scratches and stains should be inpainted while taking the larger context into consideration. Spatially homogenous filters can be used to correct unstructured faults such as film noises, image noise, and colour fading by utilising neighbouring pixels in the immediate patch.

3.1.1 Latent space translation for restoration

The restoration of historical pictures is described here as an image translation issue, where we regard fresh images and old photos as originating from separate domains and try to learn the mapping between them in order to close the domain gap. Instead of the more typical two domains, the approach translates photographs over three domains: the real photo domain, the synthetic domain, where images are intentionally damaged, and the corresponding ground truth domain Y , which includes images without degradation. This triplet domain translation is crucial to our goal since it uses both the real,

unlabeled photos and a significant quantity of artificial data that is connected to the real world.

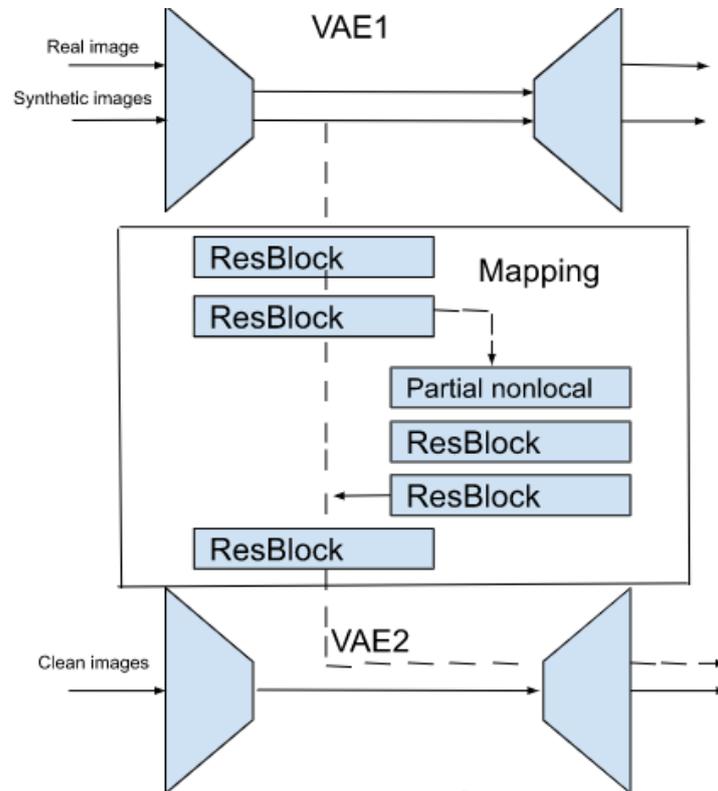


Figure 3.1: The design of our network for restoration

We primarily propose two VAEs: VAE1 for actual photographs and synthetic images, with their domain gap filled by jointly developing an adversarial discriminator; VAE2 for clear photographs. VAEs make it possible to transform images in a confined latent space. Then, we are taught the mapping that transforms corrupted images into uncorrupted ones in the latent space.

3.1.2 Variational autoencoders for image reconstruction

A sort of neural network called an autoencoder learns the data encodings from the dataset unsupervised. In essence, it consists of two parts: the first is an encoder, which resembles a convolution neural network with the exception of the last layer. The encoder's goal is to use the dataset to learn effective data encoding, which it then transfers to a bottleneck architecture. The bottleneck layer's latent space is used by the autoencoder's decoder to regenerate pictures that are similar to those in the dataset.

The loss function represents these outcomes as they backpropagate from the neural network.

In contrast to autoencoders, variational autoencoders offer a statistical method for representing the dataset's samples in latent space. The encoder generates a probability distribution in the bottleneck layer rather than a single output value as a result of variational autoencoding. Given a random vector that is created from a predefined distribution, VAE aims to produce a realistic picture. Since we did not describe the distribution of data that produces a picture, this was not feasible with the straightforward autoencoders I mentioned previously. Consequently, the plan is as follows:

- Two vectors, one for the mean and the other for the standard deviation, are produced by the encoder from a picture.
- The modified vector is the same size as the original vector after being multiplied by a tiny random number as noise. It is composed of the mean vector and the standard deviation vector.
- The decoder attempts to recreate the picture using the updated vector.
- Combining the L2 distance and the KL divergence, which quantify the departure of the distribution of the mean and standard deviation vectors from 0 and 1, we attempt to maximise the loss value.

Encourage the distribution of the mean vector to be centred around 0, whereas the distribution of the latter vector should be centred around 1. Finally, using random noise (vectors) created with a mean value of 0 and a standard deviation of 1, our decoder will be able to produce realistic pictures.

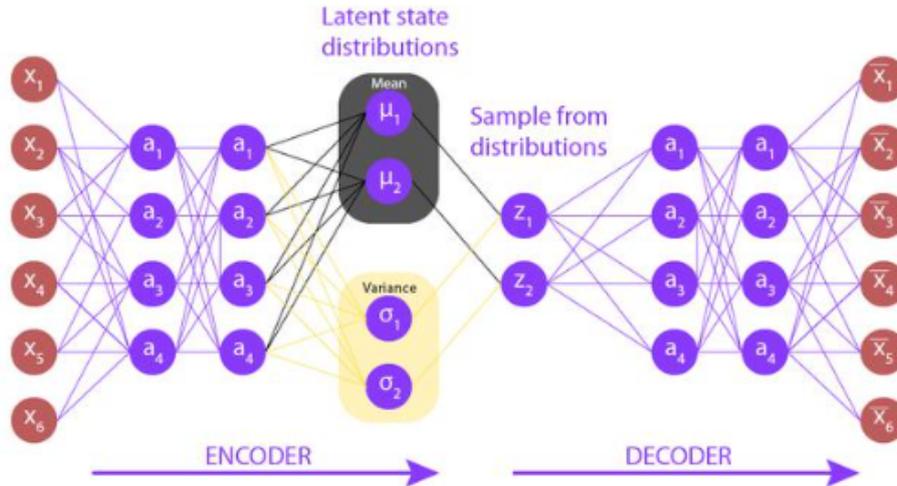


Figure 3.2: Architecture view of variational autoencoders

3.1.3 Domain arrangement in the latent space of the VAE

Our approach relies on the premise that genuine and fake pictures are encoded in the same latent spaces, and this is one of its major components. To do this, we suggest using a variational autoencoder (VAE) to encrypt pictures with a small representation, and then using an adversarial discriminator to further analyse the domain gap. To implement this idea, we utilise the network architecture depicted in the above figure.

In the first stage, two VAEs are learnt for the latent representation. The second one, VAE2, receives the encoder-generator pair together with the real ground images. The first one, known as VAE1, is shared by the encoder and generator together with old pictures and synthetic images. VAE1 is shared for both old photo and synthetic images in order to map images from both damaged domains to a single latent space. The VAEs assume a Gaussian prior for the distribution of latent codes in order to reconstruct images by sampling from the latent space. Because VAE has a denser latent representation due to KL regularisation, which minimises the domain gap and brings the latent spaces for historical photos and synthetic pictures closer with VAE1, we choose VAE over a normal autoencoder. To further reduce the domain gap in this condensed area, we propose using an adversarial network to examine the remaining latent gap.

3.1.4 Latent mapping for restoration

In the second stage, we propose to learn image restoration using synthetic picture pairings by mappings their latent space with the latent code obtained by VAEs. Latent restoration has three advantages. To begin with, because both historical and manufactured pictures are matched into the same latent space, the mapping will generalise well to recovering the photographs.

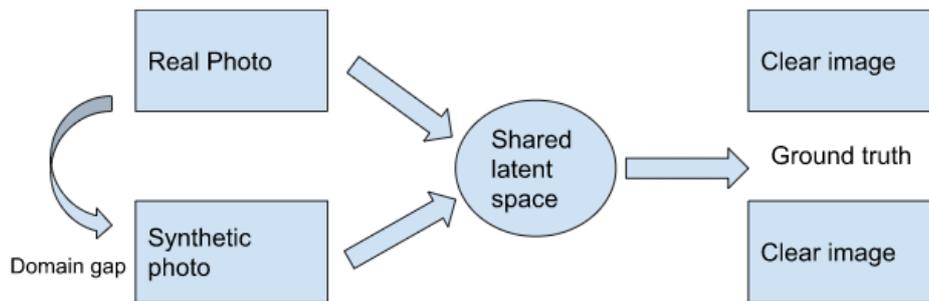


Figure 3.3: Method of restoration process

Second, compared to the high-dimensional image space, the mapping in a small low-dimensional latent space is theoretically much more simpler to learn. Additionally, since the two VAEs were trained separately, there would be no interference between the two streams when they were being rebuilt. Given the latent code transferred, the generator can always produce a perfectly clean image without any degradation; but, if we learn the translation at the pixel level, degradations will probably still exist.

3.1.5 Repeatedly restoring deterioration

The latent restoration approach stated earlier only concentrates on local features because of each layer's confined receptive field. However, realistic inpainting is required to fix structural flaws, and it must take into consideration long-range linkages to provide overall structural coherence. Since mixed degradations often appear in legacy photos, we need to develop a restoration network that can support both techniques simultaneously.

It is advised that the latent restoration network be strengthened by the inclusion of a global branch, as seen in the preceding Figure, which is made up of a nonlocal block

that takes the global context into account and several residual blocks in the following. Our nonlocal block accurately uses the mask input to ensure that the pixels in the contaminated zone will not be allowed for filling those parts, in contrast to the original block recommended in which is oblivious of the corruption region. A partial nonlocal block, which is what it is known as, is the module created particularly for the latent inpainting and is depicted in the image below. One of the features of the feature map is the context that is being considered.

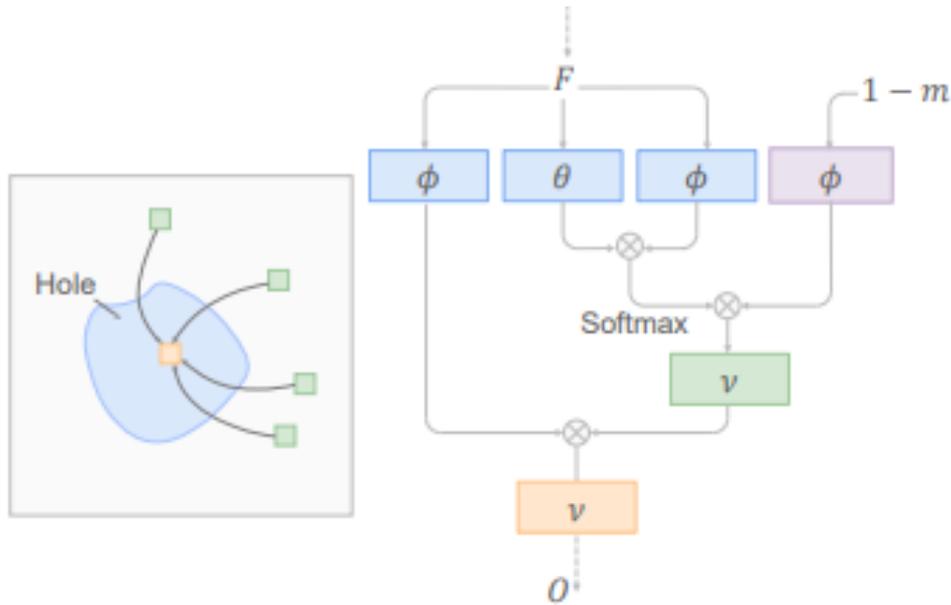


Figure 3.4: Module created particularly as a partly nonlocal block for the latent inpainting

3.1.6 Detection of Defect Regions

Because the global branch of our restoration network requires a mask as a guide, we train a supervised scratch detection network using a combination of real scratched datasets and created datasets to get the mask automatically. Given that scratch patches may occasionally make up a tiny fraction of the entire image, we employ a weight in this scenario to balance the positive and negative pixel samples.

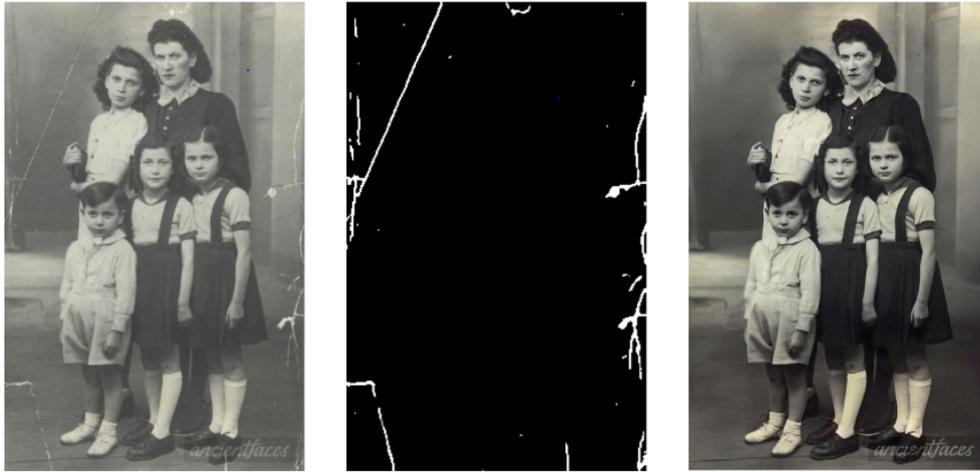


Figure 3.5: Scratch detection

3.1.7 Face Enhancement

The input picture contains faces that can be upgraded by the face enhancement model, which will then output the enhanced face image. The restoration network is applicable to all types of vintage images. However, given that human perception of restoration quality on faces is most sensitive, here we also suggest a network for face improvement. With a suggested face improvement network, we may attempt to rebuild degraded faces into a considerably more detailed and clean form given a single genuine degraded image. Due to the fact that the deterioration previous is completely unknown, the traditional pixel-wise translation approach was unable to effectively tackle this blind restoration challenge. Here, generative models to find a solution to the issue.

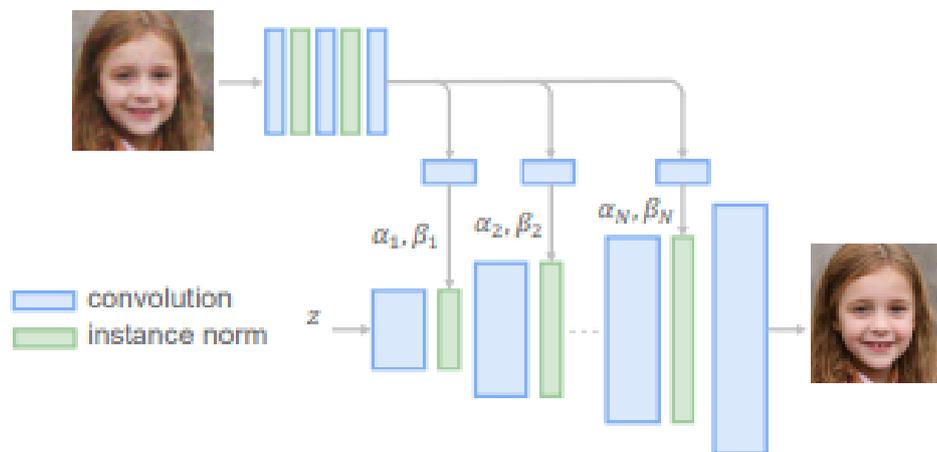


Figure 3.6: Face-enhancing progressive generator network



Figure 3.7: Examples of Face-enhancement

A network of progressive generators for face improvement is shown in the above figure. The network gradually deconvolutes the feature map starting from a latent vector. In different resolutions, the degraded face will be injected in a spatially conditional way.

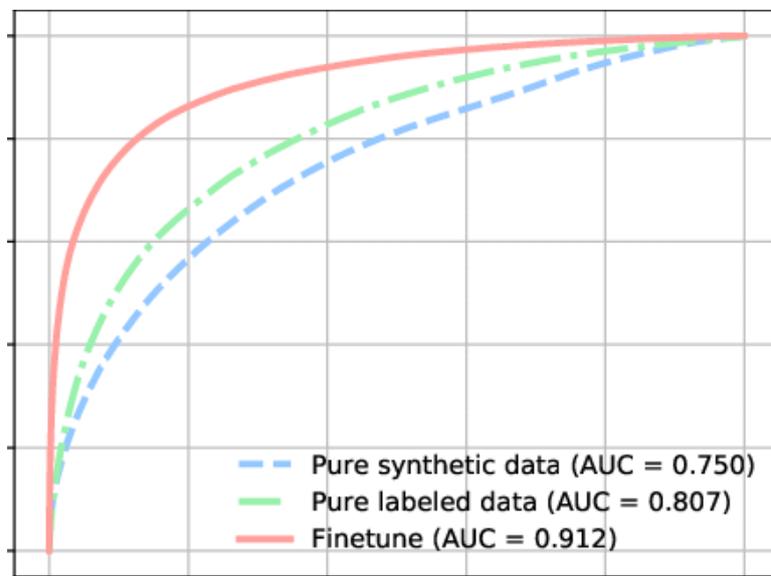


Figure 3.8: ROC curve for scratch detection

Scratch detection ROC curve for various data conditions. The scratch detection network might produce excellent results by combining synthetic organised deteriora-

tion and a little quantity of tagged data.

3.2 Dataset

The Pascal Visual Object Classes (VOC), a collection of pictures with ground truth annotation and standardised assessment tools, is made accessible to the public. Twenty object categories, including household items, domestic animals, and other, are included in the PASCAL Visual Object Classes (VOC) 2012 dataset. These include an aeroplane, a bicycle, a ship, a bus, a car, a motorbike, a train, a bottle, a chair, a plant pot, a sofa, a TV/monitor, a bird, a cat, a cattle, a dog, a horse, and a sheep. This dataset includes bounding box, object class, and pixel-level segmentation annotations for each image. For tasks including object detection, semantic segmentation, and classification, this dataset has been frequently utilised as a standard. Three subsets of the PASCAL VOC dataset were created: a private testing set, 1,449 pictures for validation, and 1,464 images for training.

- **Training Dataset:** We construct synthetic reconstructions of ancient photographs using images from the Pascal VOC collection. It also gathers scratch and paper textures, which are then refined via elastic distortions to produce realistic defects. To mix the scratch textures over the genuine images in the dataset, we utilise layer addition, lighten-only, and screen modes with varied levels of opacity. In order to mimic large-area photo damage, we make holes with feathering and haphazard forms where the underlying paper texture is visible. Finally, to replicate the unstructured faults, random amounts of film grain sounds and blurring are added.
- **Scratch detection:** We train a different network using the Unet architecture to recognise organised areas for the partial nonlocal block. The only synthetic photos are used to train the detection network at initially. To correct the imbalance of positive and negative detections, we use the focal loss. We annotate 783 old pictures that have been acquired with scratches in order to further improve the recognition performance on genuine old photos. Out of them, 400

images are used to fine-tune the detection network. The validation set's ROC curves in Figure 4 demonstrate how well finetuning works. After fine-tuning, the area under the curve (AUC) hits 0.91.

- Training details: Consider using the Adam approach in this instance, where $A1 = 0.5$ and $A2 = 0.999$. The learning rate is set at 0.0002 for the first 100 epochs, after which it progressively drops to zero. During training, we arbitrarily trim images to 256x256.

Here are some examples of synthetic photos with scratches and Some defect region detection results on real photos.



Figure 3.9: Comparison of quality against state-of-the-art techniques



Figure 3.10: synthetic photos with scratches



Figure 3.11: Some defect region detection results on real photos

3.3 Software Requirement

The software used for the project:

- Python
- Google Colab

Python

Python is a sophisticated, interactive, object-oriented, and interpreted scripting language. Python was created with the goal of being very readable. It employs more English terminology than punctuation and contains less syntactical features than other languages. It is a well-known computer programming language that is used to create apps and websites, automate activities, and analyse data. Python is a general-purpose programming language that may be used to construct a wide range of applications and is not explicitly designed to handle any specific challenges.

Data scientists and other specialists may now use Python as a standard in the field of data science to do complex statistical calculations, create machine learning techniques, handle data, and perform other tasks. The wide range of data visualisations that Python is capable of producing includes line and bar graphs, charts, histograms, and tri plots. Additionally, Python offers a number of libraries like TensorFlow and Keras that aid programmers in more quickly and efficiently developing data analysis and machine learning applications. Its features include:

- Emphasis on shorter, easier-to-read codes and readability.
- Compared to languages like C++ or Java, programmers may convey logical ideas in less code.
- Object-oriented, imperative, functional, and procedural programming are just a few of the programming paradigms that Python offers.
- Comprehensive support libraries are available.
- Dynamic language. (The given value determines the data type)

Google colab

Google is reasonably engaged in AI research. The Colaboratory platform for application development and the TensorFlow artificial intelligence framework were both developed by Google over a period of years. Google has made Colaboratory available to everyone without charge since 2017, and TensorFlow is now open-sourced. Collaboratory is currently known as Google Colab or just Colab.

The utilisation of GPU is another appealing feature that Google provides to the developers. Colab is completely free and supports GPU. Making it available to the general public for free may be done in order to establish its software as a benchmark for university courses in data science and machine learning. It could also have a long-term goal of creating a user base for the pay-per-use Google Cloud APIs. Its features includes:

- Interactive lessons for learning neural networks and machine learning.

- Without a local setup, write and run Python 3 programmes.
- Run commands on the terminal from the Notebook.
- Importing data from other sources like Kaggle.
- Notebooks can be saved to Google Drive.
- Google Drive notebooks can be imported.
- GPUs, TPUs, and cloud services for free.
- Integrate Open CV, Tensor Flow, and PyTorch.
- Easily import or publish to/from GitHub.

Chapter 4

RESULT AND DISCUSSION

Using deep learning, it is possible to repair the mixed deterioration of ancient images that have been taken in the field using a unique triplet domain translation network. The translation to clean and clear pictures is learnt in latent space, and domain gap between old photos and fake images is decreased. Comparing our strategy to earlier ones, the generalisation problem is less severe. Additionally, a partial nonlocal block for utilising the global context, recovers the latent characteristics, allowing for a better structurally consistent inpainting of the scratches.

4.1 Comparisons

All the methods were tested using the DIV2K dataset's damaged images as well as the test set of our older photo dataset. They were all trained using the identical training dataset (Pascal VOC). To allow for comparison, the following approaches are used.

- Operation-wise attention runs many processes simultaneously and chooses the best branch for mixed degradation repair using an attention technique. With supervised learning, it gains knowledge from artificial image pairings.
- Deep image prior learns the image repair from a single damaged image and has been shown to be useful in noise removal, super-resolution, and blind inpainting.
- A supervised image translation technique called Pix2Pix uses artificial picture pairings to learn the translation at the image level.

- A popular unsupervised image translation technique called CycleGAN learns the translation utilising unpaired photos from several domains.
- To repair the unstructured and structured faults, respectively, the final baseline involves performing BM3D, a traditional denoising approach, and EdgeConnect, a cutting-edge inpainting method, consecutively.

Comparing quantitatively: Four metrics are presented here to compare different models on the created images from the DIV2K dataset. The table below displays the quantitative results. When comparing the restored output to the original, the structural similarity index (SSIM) and peak signal-to-noise ratio (PSNR) are utilised to measure low-level variations.

Method	PSNR	SSIM	LPIPS	FID
Input	12.91	0.49	0.59	306.80
DIP	22.58	0.57	0.54	194.55
Attention	24.12	0.70	0.33	208.11
Pix2pix	22.18	0.62	0.23	135.14
Ours	23.33	0.69	0.25	134.35

Figure 4.1: Comparison of dataset

The user study result of various model using for image restoration process is given below.

Method	Top1	Top2	Top3	Top4	Top5
DIP	2.75	6.99	12.92	32.63	69.70
CycleGAN	3.39	8.26	15.69	24.79	52.12
Sequential	3.60	20.97	51.48	83.47	93.64
Pix2pix	14.19	54.24	72.25	86.86	96.61
Ours	64.83	81.35	90.68	96.40	98.72

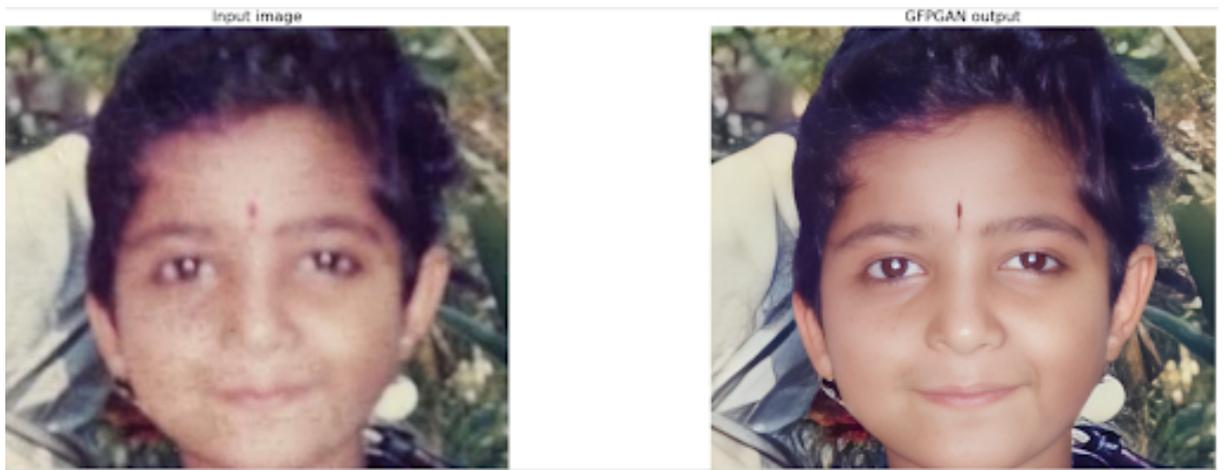
Figure 4.2: User study results



Figure 4.3: Bringing an ancient photograph back to life

4.2 Screenshots





Chapter 5

CONCLUSION

This study suggests a deep learning strategy for picture restoration and denoising. Deep learning has made incredible strides in other imaging challenges, such as image segmentation and recognition. Here, we suggest using a deep learning technique to recover vintage pictures that have suffered substantial damage. Contrary to typical restoration tasks, which can be handled by supervised learning, real-world picture deterioration is complicated, and the network is unable to generalise because to the difference in domains between synthetic images and actual old photos. Therefore, using huge amounts of synthetic picture pairings combined with real photographs, we present a unique triplet domain translation network. Comparing the approach to earlier methods, the generalisation problem is less severe. Additionally, we suggest a partial nonlocal block that, by utilising the global context, recovers the latent characteristics, allowing for a better structurally consistent inpainting of the scratches. To recreate the facial areas of ancient pictures, we also suggest a coarse-to-fine generator with spatial adaptive conditions.

When used to restore vintage images that have been extensively damaged, the approach performs well. We suggest that the characteristics of the coarse-to-fine generator be modulated in a hierarchical spatial condition way in order to rebuild a high-resolution face from genuine pictures while keeping underlying structure and style information. We contrast this strategy with a single spatial injection of several layers, ranging from the smallest scale (1616) to the largest (256256), to highlight

the significance of this issue. Qualitatively, even though we were able to produce a face that was more vivid at the smallest scale, the identity was lost since the generator could not be adequately constrained by a low-dimensional constraint. The rebuilt face steadily improves in accuracy as injection resolution is increased. However, when the injection is only done at the greatest scale, we see that the created faces have a lot of noise and artefacts. The highest scale injection point may be too near to the generator output and less important to the semantic feature in earlier layers, which causes the modulation to be insufficient.

Chapter 6

FUTURE ENHANCEMENTS

When used to restore vintage images that have been extensively damaged, the approach performs well. Complex shading, however, cannot be handled by our approach. This is because there aren't many ancient photographs in our dataset with these flaws. One might perhaps overcome this constraint by explicitly accounting for the shading effects during synthesis or by include additional photos that are comparable as training data for future advancements.

References

- [1] J. Xie, L. Xu, and E. Chen, “Image denoising and inpainting with deep neural networks,” in *Advances in neural information processing systems*, 2012, pp. 341–349.
- [2] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, “Non-local sparse models for image restoration,” in *2009 IEEE 12th international conference on computer vision*. IEEE, pp. 2272–2279.
- [3] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu, “Residual dense network for image super-resolution,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 2472–2481.
- [4] J. Sun, W. Cao, Z. Xu, and J. Ponce, “Learning a convolutional neural network for non-uniform motion blur removal,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 769–777.
- [5] S. Nah, T. Hyun Kim, and K. Mu Lee, “Deep multi-scale convolutional neural network for dynamic scene deblurring,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 3883–3891.